

NATURAL LANGUAGE RECOGNITION USING DISTRIBUTED PROCESSING

Technical Field

The present invention relates to a method of and system for natural language recognition,
 5 and in particular, to a method of and system for computer-based recognition of natural
 language data implemented on a distributed computer network.

CO-PENDING APPLICATIONS

Various methods, systems and apparatus relating to the present invention are disclosed in
 10 the following co-filed US application, the disclosures of which are incorporated herein by
 cross-reference:

NPW012 PCT

15 CROSS REFERENCES

Various methods, systems and apparatus relating to the present invention are disclosed in
 the following granted US patents and co-pending US applications filed by the applicant or
 assignee of the present application: The disclosures of all of these granted US patents and
 co-pending US applications are incorporated herein by reference.

20

	10/409,876	10/409,848	10/409,845	09/575,197	09/575,195
	09/575,159	09/575,132	09/575,123	09/575,148	09/575,130
	09/575,165	09/575,153	09/693,415	09/575,118	09/609,139
	09/608,970	09/575,116	09/575,144	09/575,139	09/575,186
25	09/575,185	09/609,039	09/663,579	09/663,599	09/607,852
	09/575,191	09/693,219	09/575,145	09/607,656	09/693,280
	09/609,132	09/693,515	09/663,701	09/575,192	09/663,640
	09/609,303	09/610,095	09/609,596	09/693,705	09/693,647
	09/721,895	09/721,894	09/607,843	09/693,690	09/607,605
30	09/608,178	09/609,553	09/609,233	09/609,149	09/608,022
	09/575,181	09/722,174	09/721,896	10/291,522	10/291,517
	10/291,523	10/291,471	10/291,470	10/291,819	10/291,481
	10/291,509	10/291,825	10/291,519	10/291,575	10/291,557

	10/291,661	10/291,558	10/291,587	10/291,818	10/291,576
	10/291,589	10/291,526	6,644,545	6,609,653	6,651,879
	10/291,555	10/291,510	19/291,592	10/291,542	10/291,820
	10/291,516	10/291,363	10/291,487	10/291,520	10/291,521
5	10/291,556	10/291,821	10/291,525	10/291,586	10/291,822
	10/291,524	10/291,553	10/291,511	10/291,585	10/291,374
	10/685,523	10/685,583	10/685,455	10/685,584	10/757,600
	09/575,193	09/575,156	09/609,232	09/607,844	09/607,657
	09/693,593	10/743,671	09/928,055	09/927,684	09/928,108
10	09/927,685	09/927,809	09/575,183	09/575,160	09/575,150
	09/575,169	6,644,642	6,502,614	6,622,999	09/575,149
	10/322,450	6,549,935	NPN004US	09/575,187	09/575,155
	6,591,884	6,439,706	09/575,196	09/575,198	09/722,148
	09/722,146	09/721,861	6,290,349	6,428,155	09/575,146
15	09/608,920	09/721,892	09/722,171	09/721,858	09/722,142
	10/171,987	10/202,021	10/291,724	10/291,512	10/291,554
	10/659,027	10/659,026	09/693,301	09/575,174	09/575,163
	09/693,216	09/693,341	09/693,473	09/722,087	09/722,141
	09/722,175	09/722,147	09/575,168	09/722,172	09/693,514
20	09/721,893	09/722,088	10/291,578	10/291,823	10/291,560
	10/291,366	10/291,503	10/291,469	10/274,817	09/575,154
	09/575,129	09/575,124	09/575,188	09/721,862	10/120,441
	10/291,577	10/291,718	10/291,719	10/291,543	10/291,494
	10/292,608	10/291,715	10/291,559	10/291,660	10/409,864
25	10/309,358	10/410,484	10/683,151	10/683,040	09/575,189
	09/575,162	09/575,172	09/575,170	09/575,171	09/575,161
	10/291,716	10/291,547	10/291,538	10/291,717	10/291,827
	10/291,548	10/291,714	10/291,544	10/291,541	10/291,584
	10/291,579	10/291,824	10/291,713	10/291,545	10/291,546
30	09/693,388	09/693,704	09/693,510	09/693,336	09/693,335
	10/181,496	10/274,119	10/309,185	10/309,066	10/778,090
	10/778,056	10/778,058	10/778,060	10/778,059	10/778,063
	10/778,062	10/778,061	10/778,057	10/782,894	10/782,895

	10/786,631	10/793,933	10/804,034	10/815,621	10/815,612
	10/815,630	HYC004US	10/815,638	10/815,640	10/815,642
	HYC008US	10/815,644	10/815,618	10/815,639	HYD001US
	10/815,647	10/815,634	10/815,632	10/815,631	10/815,648
5	10/815,614	10/815,645	10/815,646	HYG009US	10/815,620
	10/815,639	HYG012US	10/815,633	10/815,619	HYG015US
	10/815,614	10/815,636	10/815,649	10/815,609	10/815,627
	10/815,626	HYT004US	10/815,611	10/815,623	10/815,622
	HYT008US	10/815,625	10/815,624	10/815,628	10/831,232
10	10/831,242	NPS059US	NPA141US	NPT039US	NPT025US
	NPP043US	NPA150US	NPT024US	NPP040US	NPT040US
	NPT041US	NPT042US	NPT043US	NPT044US	NPK007US
	NPK006US				

- 15 Some patent applications are temporarily identified by their docket number. This will be replaced by the corresponding application number when available.

Background Art

Recent advances in pattern classification have enabled the development of sophisticated software systems that can recognize natural language data (i.e. natural language user input) such as speech (see for example L. Rabiner and B. Juang, "Fundamentals of Speech Recognition", Prentice Hall, Englewood Cliffs, NJ, 1993) or handwriting (see for example G. Lorette, "Handwriting Recognition or Reading? Situation At The Dawn of the 3rd Millennium", Advances In Handwriting Recognition, Series in Machine Perception and Artificial Intelligence, Vol. 34, pp. 3-15, World Scientific Publishing Co. 1999).

These applications allow users to communicate with a computerised system in a natural and convenient way, and permit the automation of tasks that previously required human input. Some examples of such applications include interactive voice response (IVR) systems, automated cheque-processing systems and automated form data-entry systems.

In addition, the growth of networked computing and the Internet has enabled the development of complex distributed systems, and the existence of open, standardized

protocols has allowed the integration of end-user devices, centralized servers, and applications. An example of a three-tiered distributed system architecture is depicted in Fig. 1 (prior art), illustrating a system 100 which includes a client layer 110, network layer 120 and application layer 130. Client device 140 communicates with one or more servers 150 which in turn communicate with one or more applications 160. The combination of distributed computing and pattern recognition techniques has made possible the development of systems such as NetpageTM by Silverbrook Research Pty Ltd, an interactive paper-based interface to online information. Systems such as this give users the ability to interact with information from any location that provides network connectivity (including wireless network access) using familiar human-communication techniques such as handwriting or speech.

The basic processing steps of presently known pattern recognition systems are depicted in Fig. 2 (prior art). Processing begins when an input device 210 generates a signal 220 that is to be recognized by the system 100 (that is, to be classified as belonging to a specific class or sequence of class elements). Usually, one or more pre-processing procedures 230 are applied to remove noise and produce a normalized signal 240, which is then segmented 250 to produce a stream of primitive elements 260 required for a classification procedure 270. Note that often this segmentation 250 is "soft", meaning that a number of potential segmentation points are located, and the final segmentation points are resolved during classification 270 or context processing 290.

The segmented signal 260 is then passed to a classifier 270 where a representative set of features is extracted from the signal and used in combination with a pre-defined model 275 of the input signal to produce a set of symbol hypotheses 280. These hypotheses 280 give an indication of the probability that a sequence of segments within the signal represent a basic symbolic element (e.g. letter, word, phoneme, etc.). After classification 270, the context-processing module 290 uses the symbol hypotheses 280 generated by the classifier 270 to decode the signal according to a specified context model 295 (such as a dictionary or character grammar). The result 297 produced by the context processing 290 is passed to the application 299 for interpretation and further processing.

Natural language input is inconsistent, noisy, and ambiguous, leading to potential recognition and decoding errors. However, high recognition accuracy is required for pattern recognition applications to operate successfully, since mistakes can be expensive and frustrating to users. As a result, recognition systems should make use of as much contextual information as possible to increase the possibility of correctly recognizing the natural language input. For example, when recognizing a signal that must represent a country name, the recognition system can use a pre-defined list of valid country names to guide the recognition procedure. Similarly, when recognizing a phone number, a limited symbol set (i.e. digits) can be used to constrain the recognition results. The problem domain for many pattern recognition systems is inherently ambiguous (i.e. many of the input patterns encountered during processing cannot be accurately classified without further information from a different source).

The following discussion refers to handwriting by way of background information, however, the present invention should not be considered to be limited to application to only handwriting as the form of natural language data input.

Digital ink is a digital representation of the information generated by a pen-based input device. Generally, digital ink is structured as a sequence of strokes that begin when the pen device makes contact with a drawing surface and ends when the pen-based input device is lifted. Each stroke comprises a set of sampled coordinates that define the movement of the pen-based input device whilst the pen-based input device is in contact with the drawing surface.

As an example, one of the major issues faced in the development of highly accurate handwriting recognition systems is the inherent ambiguity of handwriting (e.g. the letters 'u' and 'v', 't' and 'f', and 'g' and 'y' are often written with a very similar appearance and are thus easily confused). Human readers rely on contextual knowledge to correctly decode handwritten text, and as a result a large amount of research has been directed at applying syntactic and linguistic constraints to handwritten text recognition (see for example: H. Beigi and T. Fujisaki, "A Character Level Predictive Language Model and Its Application to Handwriting Recognition", Proceedings of the Canadian Conference on Electrical and Computer Engineering, Toronto, Canada, Sep. 13-16, 1992; U. Marti and H. Bunke,

“Handwritten Sentence Recognition”, Proceedings of the 15th International Conference on Pattern Recognition, Barcelona, Spain, Volume 3, pp. 467-470, 2000; D. Bouchaffra, V. Govindaraju, and S. Srihari, “Postprocessing of Recognized Strings Using Nonstationary Markovian Models”, IEEE Transactions Pattern Analysis and Machine Intelligence, 21(10), pp. 990-999, October 1999; J. Pitrelli and E. Ratzlaff, “Quantifying the Contribution of Language Modeling to Writer-Independent On-line Handwriting Recognition”, Proceedings of the Seventh International Workshop on Frontiers in Handwriting Recognition, Amsterdam, September 11-13 2000; R. Srihari, “Use of Lexical and Syntactic Techniques in Recognizing Handwritten Text”, ARPA Workshop on Human Language Technology, Princeton, NJ, March 1994; and L. Yaeger, B. Webb, and R. Lyon, “Combining Neural Networks and Context-Driven Search for On-Line, Printed Handwriting Recognition in the Newton”, AI Magazine, Volume 19, No. 1, pp. 73-89, AAAI 1998).

The increasing use of pen-based computing and the emergence of paper-based interfaces to networked computing resources (see for example: Anoto, “Anoto, Ericsson, and Time Manager Take Pen and Paper into the Digital Age with the Anoto Technology”, Press Release, 6th April, 2000; and Y. Chans, Z. Lei, D. Lopresti, and S. Kung, “A Feature Based Approach For Image Retrieval by Sketch”, Proceedings of SPIE Volume 3229: Multimedia Storage and Archiving Systems II, 1997) has highlighted the need for techniques to interpret digital ink. Pen-based computing allows users to interact with applications.

As a result of the progress in pen-based interface research, handwritten digital ink documents, represented by time-ordered sequences of sampled pen strokes, are becoming increasingly popular (J. Subrahmonia and T. Zimmerman: Pen Computing: Challenges and Applications. *Proceedings of the ICPR, 2000*, pp. 2060-2066). Handwriting typically involves writing in a mixture of writing styles (e.g. cursive, discrete, run-on etc.), a variety of fonts and scripts and different layouts (e.g. mixing drawings with text, various text line orientations etc.).

Presently, handwriting recognition accuracy remains relatively low, and the number of errors introduced by recognition (both for the database entries and for the handwritten query) means that present techniques do not work well. The process of converting handwriting into text results in the loss of a significant amount of information regarding the

general shape and dynamic properties of the ink. In many handwriting styles (particularly cursive writing), the identification of individual characters is highly ambiguous.

5 Similar work has been performed in the field of speech recognition, natural language processing, and machine translation.

Some known natural language recognition systems currently exist. Paragraph, Inc. offers a network-based distributed handwriting recognition system called "NetCalif" (ParaGraph, Handwriting Recognition for Internet Connected Device, November 1999) that is based on
10 their Calligraphy handwriting recognition software. The user's natural handwriting - cursive, print, or a combination of both - is captured by client software, then transmitted from an Internet-connected device to the NetCalif servers where it is converted and returned as typewritten text to the client device.

15 Philips has developed "SpeechMagic", a client/server-based, professional speech recognition software package (Philips, SpeechMagic 4.0, 2000). This system supports specialized vocabularies (called ConTexts) and dictation, recognition, and correction can be done, independently of the location, across a LAN, WAN, or the Internet.

20 In a networked information or data communications system, a user has access to one or more terminals which are capable of requesting and/or receiving information or data from local or remote information sources. The information source, in the present context, may be a database associated with an application. In such a communications system, a terminal may be a type of processing system, computer or computerised device, personal computer
25 (PC), mobile, cellular or satellite telephone, mobile data terminal, portable computer, Personal Digital Assistant (PDA), pager, thin client, or any other similar type of digital electronic device. The capability of such a terminal to request and/or receive information or data can be provided by software, hardware and/or firmware. A terminal may include or be associated with other devices, for example a pen-based input device for handwriting
30 input or a microphone for speech input.

An information source can include a server, or any type of terminal, that may be associated with one or more storage devices that are able to store information or data, such as digital

ink, for example in one or more databases residing on a storage device. The exchange of information (i.e., the request and/or receipt of information or data) between a terminal and an information source, or other terminal(s), is facilitated by a communication means. The communication means can be realised by physical cables, for example a metallic cable such as a telephone line, semi-conducting cables, electromagnetic signals, for example radio-frequency signals or infra-red signals, optical fibre cables, satellite links or any other such medium or combination thereof connected to a network infrastructure.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgment or any form of suggestion that such prior art forms part of the common general knowledge.

Disclosure Of Invention

The present invention seeks to provide improved natural language recognition, performed in a distributed system. This broadly includes a method of forwarding intermediate format data, generated by a recognizer module, to an application for context processing (i.e. decoding).

In another form, the present invention also seeks to provide means for managing multiple recognizers, user-specific dictionaries, and user-specific training of recognizers, desirable to make pattern recognition systems more accurate and flexible.

According to a first broad form of the invention, there is provided a method of providing computer-based recognition of natural language data, comprising the steps of: generating natural language data; and, transmitting the natural language data to a server; wherein, the server is programmed and configured to process the natural language data using a recognizer to produce intermediate format data, and is further capable of transmitting the intermediate format data to an application, and further wherein, the intermediate format data is decoded into computer-readable format data using context information.

According to a second broad form of the invention, there is provided a method for computer-based recognition of natural language data, comprising the steps of: receiving natural language data at a server from a remote input device; processing the natural

language data using a recognizer residing on the server to produce intermediate format data; and, transmitting the intermediate format data to an application; wherein, the application is programmed and configured to decode the intermediate format data into computer-readable format data using context information associated with the application.

5

According to a third broad form of the invention, there is provided a method of providing computer-based recognition of natural language data for interaction with an application, wherein natural language data is received at a server from a remote input device, and the server processes the natural language data using a recognizer residing on the server to
10 produce intermediate format data, the method comprising: the application receiving the intermediate format data from the server; and, the application decoding the intermediate format data into computer-readable format data using context information associated with the application.

15 According to specific, but non-limiting, embodiments of the invention, the natural language data is digital ink or speech; the digital ink is of a type from the group of: handwriting, textual, numerical, alphanumeric, pictorial or graphical; and/or the natural language data includes one or more of: normalizing the data; segmenting the data; and classifying the data.

20

According to further specific, but non-limiting, embodiments of the invention, the recognizer is implemented using software or hardware; the intermediate format data is a Directed Acyclic Graph (DAG) data structure; the DAG data structure is a matrix containing the processing results of segments of the natural language data; the intermediate
25 format data includes segmented time-series classifier data; the natural language data is derived from protein sequencing, image processing, computer vision or econometrics; the application is remote to both the input device and the server; the application resides on the server; there is more than one recognizer, each recognizer controlled by a recognition management module; the application queries the recognition management module to
30 identify a suitable recognizer to perform the processing; the context information is a user dictionary; the recognizer is able to be trained for a specific user; the input device is associated with a paper-based interface provided with coded markings; the coded markings

are a pattern of infrared markings; the input device is an optically imaging pen; and/or each paper-based interface is uniquely identified and stored on a network server.

According to a specific embodiment of the invention, there is provided a method of
5 recognising digital ink input by a user into a computer-based digital ink recognition system, the user interacting with a paper-based document, the paper-based document having disposed therein or thereon coded data indicative of a particular field of the paper-based document and of at least one reference point of the paper-based document, the method including the steps of:

10 receiving in a server, indicating data from a sensing device, operated by the user, regarding the identity of the paper-based document and at least one of a position and a movement of the sensing device relative to the paper-based document;

processing the indicating data using a recognizer residing on the server to produce intermediate format data; and,

15 transmitting the intermediate format data to an application;

wherein, the application decodes the intermediate format data into computer-readable format data using context information associated with the paper-based document;

further wherein, the sensing device comprises:

(a) an image sensor adapted to capture images of at least some of the coded data

20 when the sensing device is placed in an operative position relative to the paper-based document; and

(b) a processor adapted to:

(i) identify at least some of the coded data from one or more of the captured images;

25 (ii) decode at least some of the coded data; and

(iii) generate the indicating data using at least some of the decoded coded data.

In a particular form of the invention, the particular field of the paper-based document is
30 associated with at least one zone of the paper-based document, and the method includes identifying the context information from the at least one zone.

According to a fourth broad form of the invention, there is provided a system for computer-based recognition of natural language data, the system implemented on a network and comprising: a server to receive natural language data generated by an input device via the network; and, a recognizer residing on the server to process the natural language data to
5 produce intermediate format data; wherein, an application receives the intermediate format data and decodes the intermediate format data into computer-readable format data using context information associated with the application.

In further particular forms of the invention, the input device is a pen-based input device; the
10 input device includes a microphone; the context information is derived from one or more of a document label, a document setting, a document field label or a document field attribute; the intermediate format data is transmitted to more than one application; and/or the application initiates the processing of the natural language data.

15 According to a further aspect of the present invention there is provided a method for computer-based recognition of natural language data, the method implemented on a network and comprising the steps of:

obtaining natural language data using an input device;
receiving the natural language data on a server via the network;
20 processing the natural language data using a recognizer residing on the server to produce intermediate format data;
transmitting the intermediate format data to an application; and,
decoding the intermediate format data into computer-readable format data using context information associated with the application.

25

According to a further aspect of the present invention there is provided a method of recognising digital ink input by a user into a computer-based digital ink recognition system, the method including the steps of:

providing a user with a paper-based document, the paper-based document having
30 disposed therein or thereon coded data indicative of a particular field of the paper-based document and of at least one reference point of the paper-based document;

receiving in a server, indicating data from a sensing device, operated by the user, regarding the identity of the paper-based document and at least one of a position and a movement of the sensing device relative to the paper-based document;

5 processing the indicating data using a recognizer residing on the server to produce intermediate format data;

transmitting the intermediate format data to an application;

decoding the intermediate format data into computer-readable format data using context information associated with the paper-based document;

wherein the sensing device comprises:

10 (a) an image sensor adapted to capture images of at least some of the coded data when the sensing device is placed in an operative position relative to the paper-based document; and

(b) a processor adapted to:

15 (i) identify at least some of the coded data from one or more of the captured images;

(ii) decode at least some of the coded data; and

(iii) generate the indicating data using at least some of the decoded coded data.

20 According to a further aspect of the present invention there is provided a system for computer-based recognition of natural language data, the system implemented on a network and comprising:

an input device to generate natural language data;

a server to receive the natural language data via the network;

25 a recognizer residing on the server to process the natural language data to produce intermediate format data; and,

an application to receive the intermediate format data and to decode the intermediate format data into computer-readable format data using context information associated with the application.

30

Brief Description Of Figures

The present invention should become apparent from the following description, which is given by way of example only, of a preferred but non-limiting embodiment thereof, described in connection with the accompanying figures.

5 Fig. 1 (prior art) illustrates a distributed system architecture;

Fig. 2 (prior art) illustrates a flow chart of basic pattern recognition steps;

10 Fig. 3 illustrates an example processing system able to be used as a server to house a recognizer, according to a particular embodiment of the present invention;

Fig. 4 illustrates an example distributed recognition system, according to a particular embodiment of the present invention;

15 Fig. 5 illustrates an example of ambiguous handwriting input for "clog" / "dog";

Fig. 6 illustrates an example of ambiguous handwriting input for "tile" / "lite";

20 Fig. 7 illustrates an example recognition scenario, according to a particular embodiment of the present invention;

Fig. 8 illustrates an example recognizer selection scenario, according to a particular embodiment of the present invention;

25 Fig. 9 illustrates an example recognizer training scenario, according to a particular embodiment of the present invention;

Fig. 10 illustrates an example recognizer registration scenario, according to a particular embodiment of the present invention.

30

Modes for Carrying Out The Invention

The following modes, given by way of example only, are described in order to provide a more precise understanding of the subject matter of the present invention.

A particular embodiment of the present invention can be realised using a processing system, an example of which is shown in Fig. 3. In particular, the processing system 300 generally includes at least one processor 302, or processing unit or plurality of processors, memory 304 and at least one output device 308, coupled together via a bus or group of buses 310. At least one storage device 314 which houses at least one database 316 can also be provided, which may be remote and accessed via a network. The memory 304 can be any form of memory device, for example, volatile or non-volatile memory, solid state storage devices, magnetic devices, etc. The processor 302 could include more than one distinct processing device, for example to handle different functions within the processing system 300.

Input device 306, for example a pen-based input device or a microphone, is normally remote to the system 300. Input device 306 is used by a user to generate natural language data 318 which is preferably transmitted over network 307 to system 300 for processing. Output device 308 produces or generates intermediate format data 320, for example for transmission over a network, to be transmitted to application 324, which could be remote or local to the system 300. The storage device 314 can be any form of data or information storage means, for example, volatile or non-volatile memory, solid state storage devices, magnetic devices, etc.

In use, the processing system 300 may be a server and is adapted to allow data or information to be stored in and/or retrieved from, via wired or wireless communication means, the at least one database 316, which may be remote and accessed via a further network. The processor 302 receives natural language data 318 from input device 306, preferably via network 307, and outputs intermediate format data 320 by utilising output device 308, for example a network interface. The application 324 may return decoded data to the processing system. The application 324 may cause information to be printed, for example on a NetpageTM printer, at a user's location. More than one input device 306 can be provided. It should be appreciated that the processing system 300 may be any form of terminal, server, specialised hardware, or the like. The processing system 300 may be a part of a networked communications system. Also, the application 324 may initiate transfer of natural language data 318 from the input device 306 to server 300.

In a particular embodiment, the server 300 is part of a system for computer-based recognition of natural language data, the system implemented on a network and comprising: the input device 306 to obtain natural language data; server 300 to receive the natural language data 318 via a network 307; a recognizer residing on the server 300 to process, in processor 302, the natural language data 318 to produce intermediate format data 320; and, an application 324 to receive the intermediate format data 320 and to decode the intermediate format data 320 into computer-readable format data using context information associated with the application 324.

The following example provides a more detailed discussion of a particular embodiment of the present invention. The example is intended to be merely illustrative and not limiting to the scope of the present invention.

In a particular preferred embodiment, the present invention is configured to work with the NetpageTM networked computer system, a detailed description of which is given in the applicant's co-pending applications, including in particular, PCT Publication No. WO0242989 entitled "Sensing Device" filed 30 May 2002, PCT Publication No. WO0242894 entitled "Interactive Printer" filed 30 May 2002, PCT Publication No. WO0214075 "Interface Surface Printer Using Invisible Ink" filed 21 February 2002, PCT Publication No. WO0242950 "Apparatus For Interaction With A Network Computer System" filed 30 May 2002, and PCT Publication No. WO03034276 entitled "Digital Ink Database Searching Using Handwriting Feature Synthesis" filed 24 April 2003.

It will be appreciated that not every implementation will necessarily embody all or even most of the specific details and extensions described in these applications in relation to the basic system. However, the system is described in its most complete form to assist in understanding the context in which the preferred embodiments and aspects of the present invention operate.

In brief summary, the preferred form of the Netpage system provides an interactive paper-based interface to online information by utilizing pages of invisibly coded paper and an optically imaging pen. Each page generated by the Netpage system is uniquely identified

and stored on a network server, and all user interaction with the paper using the Netpage pen is captured, interpreted, and stored. Digital printing technology facilitates the on-demand printing of Netpage documents, allowing interactive applications to be developed. The Netpage printer, pen, and network infrastructure provide a paper-based alternative to
5 traditional screen-based applications and online publishing services, and supports user-interface functionality such as hypertext navigation and form input.

Typically, a printer receives a document from a publisher or application provider via a broadband connection, which is printed with an invisible pattern of infrared tags that each
10 encodes the location of the tag on the page and a unique page identifier. As a user writes on the page, the imaging pen decodes these tags and converts the motion of the pen into digital ink. The digital ink is transmitted over a wireless channel to a relay base station, and then sent to the network for processing and storage. The system uses a stored description of the page to interpret the digital ink, and performs the requested actions by interacting with an
15 application.

Applications provide content to the user by publishing documents, and process the digital ink interactions submitted by the user. Typically, an application generates one or more interactive pages in response to user input, which are transmitted to the network to be
20 stored, rendered, and finally printed as output to the user. The Netpage system allows sophisticated applications to be developed by providing services for document publishing, rendering, and delivery, authenticated transactions and secure payments, handwriting recognition and digital ink searching, and user validation using biometric techniques such as signature verification.

25

Distributed Pattern Recognition

An example architecture for a distributed pattern recognition system 400 is depicted in Fig. 4. In the example, a signal 410 is recorded by an input device 415 at a client layer 420 and transmitted over a network to a server (network layer 430) for recognition by a recognizer
30 440, with the intermediate results 445 transmitted back to the client layer 420 or a third party application 450 on an application layer 455 for interpretation and processing. One advantage of this approach is that client devices 415 and distributed applications 450 do not require the significant computing resources commonly needed to perform natural language

pattern recognition, and the network servers that perform the recognition are not subject to the resource constraints that are inherent in many client devices 415 (e.g. mobile phones, personal-digital assistants, imaging pens, etc.). As a result, network servers are able to use extremely processor- and/or memory-intensive techniques to improve recognition accuracy, and can use hardware optimised to perform the specific recognition task.

Performing pattern recognition on a centralized server (e.g. processing system 300) also offers an advantage to pattern-recognition systems that employ user-specific adaptation to achieve higher recognition rates. For example, some handwriting recognition techniques develop a handwriting model for each user of the system based on previous recognition results, which is then used to improve the future accuracy of the system for that user (see for example L. Schomaker, H. Teulings, E. Helsper, and G. Abbink, "Adaptive Recognition Of Online, Cursive Handwriting", Proceedings of the Sixth International Conference on Handwriting and Drawing. Paris, July, 4-7 Telecom, (pp. 19-21), 1993 and S. Connell and A.K. Jain, "Writer Adaptation of Online Handwritten Models," Proc. 5th International Conference on Document Analysis and Recognition, Bangalore, India, pp. 434-437, Sept. 1999).

This adaptation is more effective if a single server, or set of servers, performs all recognition for a user (rather than a large number of individual applications each performing their own recognition), since the server is able to perform adaptation based on the input generated by all applications. In addition to this, centralized server-based pattern recognition simplifies the management of the recognition system 400 by allowing recognizers to be reconfigured and upgraded without interaction with the distributed client devices 415 and applications 450, and allows training and test data to be easily collected.

However, the information required to perform the context processing stage of a pattern recognition system is generally application specific and is often very large (e.g. entries in a large application-specific database), making it impractical to transmit the context information to a centralized server for processing. A solution to this problem is to use a mechanism for distributed recognition as depicted in Fig. 4. When a user generates a signal (i.e. natural language data) 410 to be recognized and processed by an application, the signal 410 is submitted to a distributed server for processing. The server performs processing

steps such as pre-processing, segmentation, and classification (see Fig. 2), but does not use a context model to decode the result (or only performs partial decoding as described in the following discussion). Rather, the intermediate recognition results (i.e. intermediate format data) are returned or sent to the application allowing the application to apply any arbitrarily
5 complex and domain-specific context processing to decode the signal.

Symbol DAG

One method of returning the intermediate recognition results (i.e. intermediate format data) to an application is to use a symbol DAG (Directed Acyclic Graph), which is a generic data
10 structure that contains symbol and associated scores as vertices, and valid transitions between symbols as edges. The structure can be implemented as a two-dimensional array of elements, each of which defines the output generated by the pattern classifier for a single segment of the signal and the associated valid transitions for that segment. This structure represents all the potential recognition alternatives that may be derived from the input
15 signal based on the results of the classifier. The application uses this structure, in combination with a context model, to decode the input signal.

The symbol DAG is equivalent to a matrix where each column contains the results of the classification of a single segment of the input signal. Each element in the column represents
20 the probability that the classified segment is a particular symbol, and includes an offset that indicates the next possible segment (column) in the input signal that can follow this symbol. Thus, the matrix represents all the possible decoding paths based on the output of the pattern classifier. These paths and associated classification scores can be combined with a context model to fully decode the input signal.

25

Note that the symbol DAG is applicable in any pattern recognition task where a sequence of classification results is decoded using a context or set of constraints. The symbols contained in the symbol DAG may be any primitive element that is generated as the output of a pattern classifier, including the output from a time-series classifier. Examples of such
30 recognition systems include handwriting and speech recognition, protein sequencing (see A. C. Camproux, P. Tuffery, S. Hazout, "Hidden Markov Model Approach For Identifying The Modular Framework Of The Protein Backbone", Protein engineering, 12(12), pp. 1063, December 1999), image processing and computer vision (see Y. He, A. Kundu, "2-D Shape

Classification Using Hidden Markov Model", IEEE Transactions on Pattern Analysis, 13(11), November 1991), and econometrics (see T. Ryden, T. Terasvirta, S. Asbrink, "Stylized Facts of Daily Return Series and the Hidden Markov Model", Journal of Applied Econometrics, 13(3), pp. 217, May 1998).

5

Symbol DAG Example

As an example, Table 1 shows a symbol DAG that represents the output from a handwritten character recognizer generated by the ambiguous text given in Fig. 5. In this example, the recognizer has found two possible character segmentation arrangements, as depicted by the
 10 two rows in the symbol DAG. Note that in the examples, the symbol scores are given as probabilities; however, an actual implementation may typically use log-probabilities (i.e. the base-10 logarithm of the probability result) to improve the performance of context processing and to avoid overflow and underflow problems that occur when multiplying probabilities using finite precision floating-point operations.

15

To decode the alternatives, the context processor starts with the first entry in the DAG (i.e. the character 'c'). The score for this entry is added to the accumulated total (since log-probabilities are added rather than multiplied), and processing moves to the column given by the offset value in the entry (in this example, column 1). In column 1, two alternatives
 20 exist (i.e. "cl" or "cb"), and the scores for these alternatives are found by adding the scores to the previous total. The decoding continues until the end of the DAG is reached. Similarly, the second entry in column 0 (i.e. the character 'd') is decoded; note however, that column 1 is skipped in this traversal of the DAG, as indicated by the offset value of 2 in the character score entry. This is due to the letter 'd' being constructed using two strokes,
 25 and thus the recognition of the letters 'l' and 'b' cannot be valid in this alternative. Thus, the potential decoding alternatives in this example are:

$$\text{clog} = 0.7 * 0.8 * 1.0 * 1.0 = 0.56$$

30

$$\text{cbg} = 0.7 * 0.2 * 1.0 = 0.14$$

$$\text{dog} = 0.3 * 1.0 * 1.0 = 0.30$$

These values can now be combined with a language model or other contextual information to select the most likely word.

	0	1	2	3
Character	c	l	o	g
Offset	1	2	3	0
Score	0.7	0.8	1.0	1.0
Character	d	b		
Offset	2	3		
Score	0.3	0.2		

5

Table 1. Example DAG for "clog"/"dog" ambiguity

The DAG structure must ensure that strokes are assigned to an individual letter only once. To do this, alternate paths must be defined to ensure that if a stroke is assigned to a letter, no subsequent letter may use that stroke in its construction. An example of this is given in Fig. 6, with the derived DAG depicted in Table 2. In this example, the short, horizontal marks can potentially be recognized as crossbar elements of a letter 't', or diacritical marks for the letter 'i'. However, if a marking is used as a crossbar, it cannot subsequently be used as a diacritical. The potential decoding alternatives in this example are:

10

15

$$\text{tile} = 0.6 * 1.0 * 0.6 * 1.0 = 0.36$$

$$\text{tite} = 0.6 * 1.0 * 1.0 * 1.0 = 0.60$$

$$\text{lite} = 0.4 * 1.0 * 1.0 * 1.0 = 0.40$$

20

These values can now be combined with a language model to select the most likely word.

	0	1	2	3	4	5
Character	t	i	i	t	l	e
Offset	1	4	3	5	5	-
Score	0.6	1.0	1.0	1.0	0.6	1.0

Character	1		t
Offset	2		5
Score	0.4		0.4

Table 2. Example DAG for "lite"/"tile" ambiguity

5 Additionally, the character value of a DAG entry can be set to zero, indicating a NUL character (i.e. a character that does not change the text, but will modify the text probability). This allows word break positions (i.e. spaces) to be modeled as a SPACE/NUL pair, indicating that there is a certain probability that a space appears at that point in the DAG. For example:

	0	1	2
Character	a	NUL	b
Offset	1	1	-
Score	1.0	0.6	1.0
Character		SPACE	
Offset		1	
Score		0.4	

Table 3. Example DAG for SPACE/NUL pair

5 The potential decoding alternatives in this example are:

$$ab = 1.0 * 0.6 * 1.0 = 0.6$$

$$a b = 1.0 * 0.4 * 1.0 = 0.4$$

10

Distributed Recognizer Management

Referring to Figs. 7 and 8, a distributed recognition system 700 may support a number of different recognizers 440 that are controlled by a distributed recognition management system or recognition manager 710. These recognizers 440 can include systems capable of supporting different classes of recognition, such as different languages, dialects, or accents, or cursive or boxed input for handwriting systems. When an application 450 requires a recognition task to be performed, the application 450 first queries 720 the recognition manager 710 to find a recognizer 440 that matches the parameters of the input to be recognized (as depicted in Fig. 8). The recognition manager 710 then queries 730 each recognizer 440 to find a recognizer that supports the parameters specified by the application 450. When a recognizer 440 indicates support 740 (as opposed to no support 750 from recognizer 440a in Fig. 8) for the specified parameter set, the enumeration ends and the selected recognizer 440 (in the case of Fig. 8 recognizer 440b) is passed 760 to the application 450. Note that the individual recognizers 440 do not need to be centralized and may be distributed throughout the system 700, since the recognition manager 710 acts as a controller for the set of recognizers 440. The application 450 can then request processing by

the selected recognizer by passing or directing 770 the signal and parameters to the selected recognizer 440. Intermediate format data 445, i.e. a symbol lattice, is returned to the application 450 and the application 450 can return a response 780 to the input device 415.

5 User-Specific Dictionaries

Distributed recognition systems can also support user dictionaries, which are user-specific word lists (and possibly associated a-priori probabilities) that include words that a user writes frequently but which are unlikely to appear in a standard dictionary (examples include company names, work or personal interest specific terms, etc.). User dictionaries
10 can be stored and managed centrally so that words added to the dictionary when using one application are available to all applications for context processing. Obviously, applications can manage and use their own local user-specific dictionaries if required, since they have full control over context decoding.

15 When an application requires the recognition of a signal that may contain words found in the user dictionary (e.g. standard handwritten text input such as the subject line of an e-mail or an arbitrary voice message), the centralized recognition system generates the usual intermediate recognition results to be returned to the application for context decoding. However, in addition to this it decodes the intermediate results using the user-dictionary as
20 a language model, the result of which is also returned to the application. These two intermediate results structures can be combined by the application during its context decoding to generate a final decoding that includes the user-specific dictionary information.

User-Specific Training

25 Distributed recognition systems may also support user-specific training for a recognizer 440, as depicted in Fig. 9. The data generated by a user-specific recognition training application is submitted 910 to the centralized recognition manager 710, which stores 920 the data in a database 930. The recognition manager 710 then enumerates all recognizers 440 to determine if they support the data format as defined by the parameters associated
30 with the training data, and if so (True signal 940), submits the training data 950 to the recognizer 440 for user-specific training.

When an existing recognizer is upgraded or a new recognizer is added to the system, the recognition manager 710 queries 1010 the training database 930 to determine if any training data 1020 of the format required by the recognizer 440 exists. If so, the training data 1020 is submitted to the newly registered recognizer 440 for processing, as depicted in Fig. 10.

5

The invention may also be said to broadly consist in the parts, elements and features referred to or indicated herein, individually or collectively, in any or all combinations of two or more of the parts, elements or features, and wherein specific integers are mentioned herein which have known equivalents in the art to which the invention relates, such known

10 equivalents are deemed to be incorporated herein as if individually set forth.

Although a preferred embodiment has been described in detail, it should be understood that various changes, substitutions, and alterations can be made by one of ordinary skill in the art without departing from the scope of the present invention.

15